Streaming View Learning

Chang Xu xuchang@pku.edu.cn
Key Laboratory of Machine Perception (Ministry of Education)
School of Electronics Engineering and Computer Science
Peking University
Beijing 100871, China

Dacheng Tao dacheng.tao@uts.edu.au
Centre for Quantum Computation and Intelligent Systems
Faculty of Engineering and Information Technology
University of Technology, Sydney
Sydney, NSW 2007, Australia

Chao Xu xuchao@cis.pku.edu.cn
Key Laboratory of Machine Perception (Ministry of Education)
School of Electronics Engineering and Computer Science
Peking University
Beijing 100871, China

arXiv:1604.08291v1 [stat.ML] 28 Apr 2016

Abstract

An underlying assumption in conventional multi-view learning algorithms is that all views can be simultaneously accessed. However, due to various factors when collecting and pre-processing data from different views, the streaming view setting, in which views arrive in a streaming manner, is becoming more common. By assuming that the subspaces of a multi-view model trained over past views are stable, here we fine tune their combination weights such that the well-trained multi-view model is compatible with new views. This largely overcomes the burden of learning new view functions and updating past view functions. We theoretically examine convergence issues and the influence of streaming views in the proposed algorithm. Experimental results on real-world datasets suggest that studying the streaming views problem in multi-view learning is significant and that the proposed algorithm can effectively handle streaming views in different applications.

Keywords: Multi-view Learning, Streaming Views

1. Introduction

In this era of exponential information growth, it is now possible to collect abundant data from different sources to perform diverse tasks, such as social computing, environmental analysis, and disease prediction. These data are usually heterogeneous and possess distinct physical properties such that they can be categorized into different groups, each of which is then regarded as a particular view in multi-view learning. For example, in video surveillance (Wang, 2013), placing multiple cameras at different positions around one area might enable better surveillance of that area in terms of accuracy and reliability. In another example, to accurately recommend products to target customers (Jin et al., 2005), it is necessary to comprehensively describe the product by its image, brand, supplier, sales history, and user

feedback. Effective descriptors have already been developed for object motion recognition: (i) histograms of oriented gradients (HOG) (Dalal and Triggs, 2005), which focus on static appearance information; (ii) histograms of optical flow (HOF) (Laptev et al., 2008), which capture absolute motion information; and (iii) motion boundary histograms (MBH) (Dalal et al., 2006), which encode related motion between pixels.

Rather than requiring that all the examples should be comprehensively described based on each individual view, it might be better to exploit the connections and differences between multiple views to better represent examples. A number of multi-view learning algorithms (Blum and Mitchell, 1998; Lanckriet et al., 2004; Jia et al., 2010; Chen et al., 2012; Chaudhuri et al., 2009) have thus emerged to effectively fuse multiple features from different views. These have been widely applied to various computer vision and intelligent system problems.

Although an optimal description of the data might be obtained by integration of multiple views, in practice it is difficult to guarantee that all the candidate views can be accessed simultaneously. For example, establishing a camera network for video surveillance is a huge project that takes time to realize. The number of views used in tracking and detection has increased. Newly developed recommendation systems might well have their images and text descriptions in place, but they require a period of time to accumulate sales and, therefore, user feedback, which are key factors influencing the decisions of prospective customers. Images can be depicted by diverse visual features with distinct acquisition costs. For example, a few milliseconds might be sufficient to extract the color histogram or SIFT descriptors from a normal-sized image, but time-cost clustering and mapping processes are further required to generate bag-of-word (BoW) features from SIFT descriptors. Recent deep learning methods need longer time (usually hours or even days) to obtain a reasonable model for image feature extraction.

Conventional multi-view learning algorithms (Cai et al., 2013; Kumar and Udupa, 2011) have been developed in ideal settings in which all the views are accessed simultaneously. The real world, however, presents a more challenging multi-view learning scenario formed from multiple streaming views. Newly arrived views might contain fresh and timely information, that are beneficial for further improving the multi-view learning performance. To make existing multi-view learning methods applicable to this streaming view setting, a naive approach might be to treat a new view arriving as a new stage each time and then running the multi-view learning algorithms again with the new views. However, this approach is likely to suffer from intensive computational costs or serious performance degradation. In contrast, here we propose an effective streaming view learning algorithm that assumes the view function subspaces in the well-trained multi-view model over sufficient past views are stable and fine tunes their combination weights for an efficient model update. We provide theoretical analyses to support the feasibility of the proposed algorithm in terms of convergence and estimation error. Experimental results in real-world clustering and classification applications demonstrate the practical significance of investigating streaming views in the multi-view learning problem and the effectiveness of the proposed algorithm.

2. Problem Formulation

In the standard multi-view learning setting, we are provided with n examples of m views $\{(x_i^1, \cdots, x_i^m)\}_{i=1}^n$, where $x_i^v \in \mathbb{R}^{D_v}$ is the feature vector on the v-th view of the i-th example.

The feature matrices on different views are thus denoted as $\{X^1, \cdots, X^m\}$, where $X^v \in \mathbb{R}^{D_v \times n}$. Subspace-based multi-view learning approaches aim to discover a subspace shared by multiple views, such that the information from multiple views can be integrated in that subspace:

$$(f^1, \cdots, f^m) : (x_i^1, \cdots, x_i^m) \to z_i, \tag{1}$$

where f^v is the view function on the v-th view, and $z_i \in \mathbb{R}^d$ is the latent representation in the subspace $\mathcal{Z}$. Based on the unified representation z of the multi-view example, the subsequent tasks, including classification, regression, and clustering, can easily be accomplished.

Most existing multi-view learning algorithms explicitly assume that all views $\{X^1, \cdots, X^m\}$ are static and can be simultaneously accessed for multi-view model learning. If new views $\{X^{m+1}, \cdots, X^{m+k}\}$ are provided, the question arises of how to upgrade the well-trained multi-view model $(f^1, \cdots, f^m)$ over the past m views using the latest information. It is unreasonable to simply neglect the newly arrived views and ignore the possibility of updating the model. On the other hand, naively maintaining a training pool composed of all views, enriching the pool with each newly arrived view, and then re-launching the multi-view learning algorithm will be resource (storage and computation) consuming. It is, therefore, necessary to investigate this challenging multi-view learning problem in the general streaming view setting, where multiple views arrive in a streaming format.

2.1 Streaming View Learning

In this section, we first present a naive approach to handle new views. We then develop a sophisticated streaming view learning algorithm that reduces the burden of learning new views.

2.1.1 A Naive Approach

Assume that multiple example views $\{x^1, \cdots, x^m\}$ are generated from a latent data point z in the subspace,

$$x^v = f^v(z) = W_v z, \tag{2}$$

where view function f^v parameterized by $W_v \in \mathbb{R}^{D_v \times d}$ can be assumed to be linear for simplicity. In practice, different feature dimensions are usually correlated; for example, distinct image tags in BoW features might be related to each other. It is thus reasonable to encourage a low rank of W_v. Moreover, the low-rank $\{W_v\}_{v=1}^m$ implies that the latent subspace contains comprehensive information to generate multiple view spaces while the inverse procedure is infeasible, which is consistent with our assumption that multiple views are generated from a latent subspace.

Within the framework of empirical risk minimization, view functions $W = \{W_1, \cdots, W_m\}$ can be solved with the following problem:

$$\min_{W,z} \frac{1}{nm} \sum_{i=1}^{n} \sum_{v=1}^{m} \|x_i^v - W_v z_i\|_2^2 + C_1 \sum_{v=1}^{m} \|W_v\|_* + C_2 \sum_{i=1}^{n} \|z_i\|_2^2, \tag{3}$$

where a least-squared loss is employed to measure the reconstruction error of multi-view examples, and a trace norm is applied to regularize view functions.

Suppose that, by solving Problem (3), we have already a well-trained multi-view model $\{W_1, \cdots, W_m\}$ over m views $\{X^1. \cdots, X^m\}$. Considering a new arriving view X^{m+1}, we

are then faced with challenging problems of how to discover the view function W_{m+1} for the new view and how to upgrade the view functions $\{W_1, \cdots, W_m\}$ on past m views. It is a straightforward extension to simultaneously handle more than one new view.

Within the framework of Eq. (3), the latent representations $\{z_i\}_{i=1}^n$ previously learned over m past views are regarded as fixed. Then, the new view function W_{m+1} can be efficiently solved by

$$\min_{W_{m+1}} \frac{1}{n} \sum_{i=1}^{n} \|x_i^{m+1} - W_{m+1} z_i\|_2^2 + C_1 \|W_{m+1}\|_*. \tag{4}$$

Problem (3) can then be naturally adapted to $m+1$ views, and alternately optimizing view functions $\{W_v\}_{v=1}^{m+1}$ and subspace representations $\{z_i\}_{i=1}^n$ for several iterations will output the optimal multi-view model.

This naive approach to handling streaming views can be treated as a stochastic optimization strategy for solving Problem (3). The view function for the new view can be efficiently discovered with the help of latent representations learned on past views; however, it is computationally expensive to upgrade the view functions on past views by re-launching Problem (3), especially when the number of views and the view function dimensions are large.

2.1.2 Streaming View Learning

We begin the development of the streaming view learning algorithm by carefully investigating the view function. Note that any matrix $W_v \in \mathbb{R}^{D_v \times d}$ can be represented as the sum of rake-one matrices:

$$W_v = \sum_{ij} \sigma_{ij}^v a_i^v (b_j^v)^T, \tag{5}$$

where $span(a_1^v, a_2^v, \cdots) = \mathbb{R}^{D_v}$ and $span(b_1^v, b_2^v, \cdots) = \mathbb{R}^d$, and $\{\sigma_{ij}^v\}$ are the coefficients to combine different subspaces.

Based on the new formulation of the view function in Eq. (5), Problem (3) can be reformulated as

$$\begin{aligned} &\min \frac{1}{nm} \sum_{i=1}^{n} \sum_{v=1}^{m} \|x_i^v - A_v S_v B_v^T z_i\|_2^2 + C_1 \sum_{v=1}^{m} \|S_v\|_* + C_2 \sum_{i=1}^{n} \|z_i\|_2^2 \\ &\text{w.r.t. } \forall v\ A_v \in \mathbb{R}^{D_v \times k_v}, B_v \in \mathbb{R}^{d \times k_v}, S_v \in \mathbb{R}^{k_v \times k_v}; \forall i\ z_i \in \mathbb{R}^d \\ &\text{s.t. } \forall v\ A_v^T A_v = I, \quad B_v^T B_v = I, \end{aligned} \tag{6}$$

where A_v and B_v correspond to the column and row spaces of W_v respectively, S_v contains the weights to combine different rank-one subspaces, and k_v indicates the number of active function subspaces on the v-th view.

Suppose that we already have well trained view functions $\{(A_v, B_v, S_v)\}_{v=1}^m$ over m views. For the new $(m+1)$-th view, we can efficiently discover its view function $(A_{m+1}, B_{m+1}, S_{m+1})$

given the fixed latent representations $\{z_i\}_{i=1}^n$,

$$\begin{aligned} \min\ & \frac{1}{n}\sum_{i=1}^{n} \|x_i^{m+1} - A_{m+1}S_{m+1}B_{m+1}^T z_i\|_2^2 + C_1\|S_{m+1}\|_* \\ \text{w.r.t.}\ & A_{m+1} \in \mathbb{R}^{D_{m+1}\times k_{m+1}},\ B_{m+1} \in \mathbb{R}^{d\times k_{m+1}}, \\ & S_{m+1} \in \mathbb{R}^{k_{m+1}\times k_{m+1}} \\ \text{s.t.}\ & A_{m+1}^T A_{m+1} = I, \quad B_{m+1}^T B_{m+1} = I. \end{aligned} \tag{7}$$

The remaining task is to then upgrade the view functions on past m views using the latest information. As mentioned above, completely re-training the model on past views is computationally expensive since a large number of variables need to be learned. Instead, we propose to fine-tune the previously well-trained multi-view model using the following objective function:

$$\begin{aligned} \min\ & \frac{1}{n(m+1)}\sum_{i=1}^{n}\sum_{v=1}^{m+1} \|x_i^v - A_v S_v B_v^T z_i\|_2^2 \\ & + C_1\sum_{v=1}^{m}\|S_v\|_* + C_2\sum_{i=1}^{n}\|z_i\|_2^2 \\ \text{w.r.t.}\ & \forall v\ S_v \in \mathbb{R}^{k_v\times k_v}; \quad \forall i\ z_i \in \mathbb{R}^d, \end{aligned} \tag{8}$$

where we have fixed the row and column spaces of view functions on multiple views and attempted to update view functions by adjusting their coefficients for subspace combination. Since the view functions are now mainly determined by a set of smaller matrices $\{S_v\}_{v=1}^{m+1}$, where $S_v \in \mathbb{R}^{k_v\times k_v}$ with $k_v \ll \min(D_v, d)$, solving Problem (8) is often much cheaper than solving Problem (6) or (3) with $m+1$ views.

After solving or updating the view functions on $(m+1)$ views, the multi-view model can then process another new view. Meanwhile, the current multi-view model can be used to predict the latent representation of a new multi-view example followed by subsequent tasks.

3. Optimization

The proposed streaming view learning algorithm involves optimization over latent representations $z = \{z_i\}_{i=1}^m$ and function subspaces on multiple views $\{(A_v, B_v)\}_{v=1}^m$ and their corresponding combination weights $\{S_v\}_{v=1}^m$. In this section, we employ an alternating minimization strategy to optimize these variables. The whole optimization procedure is summarized in Algorithm 1.

3.1 Optimization Over Latent Representations

Fixing view functions $\{W_v = A_v S_v B_v^T\}_{v=1}^m$ on multiple views, the optimization problem w.r.t. the latent representation of the i-th example is

$$\min_{z_i}\ \frac{1}{nm}\sum_{v=1}^{m}\|x_i^v - W_v z_i\|_2^2 + C_2\|z_i\|_2^2, \tag{9}$$

which is easy to solve in a closed form.

3.2 Optimization Over View Function Subspaces

By fixing the latent representations, the view functions on multiple views can be independently optimized via

$$\min_{W_v} \ g(W_v) + C_1\|W_v\|_*, \tag{10}$$

where

$$g(W_v) = \frac{1}{nm}\sum_{i=1}^{n}\|x_i^v - W_v z_i\|_2^2. \tag{11}$$

The proximal gradient descent method (Ji and Ye, 2009) has been widely used to solve this problem by reformulating it to,

$$\min_{W_v} \ \frac{\eta_t}{2}\|W_v - \big(W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1})\big)\|_F^2 + C_1\|W_v\|_*, \tag{12}$$

where η_t is the step size in the t-th iteration. It turns out that Problem (12) can be solved by singular value thresholding (SVT) (Cai et al., 2010),

$$W_v^t = soft\big(W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1}), \frac{C_1}{\eta_t}\big), \tag{13}$$

where $soft(X, C) = A(\Sigma - CI)_+B^T$ with singular value decomposition $X = A\Sigma B^T$ for X.

By operating Eq. (13), we can obtain the view function subspace. However, Eq. (13) requires accurate SVD over $W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1})$, which is computationally expensive given the large dimension of W_v. Recall that this step is only used to identify the view function subspaces, and the view function is more accurately discovered by optimizing the combination weight. Therefore, it is unnecessary to compute the SVD of $W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1})$ very accurately. We apply the power method (Halko et al., 2011) with several iterations to approximately calculate $W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1}) \approx \widehat{A}_v\widehat{\Sigma}_v\widehat{B}_v^T$.

We initialize the optimization method by $(W_v)_0 = X_v Z^T(ZZ^T)^{-1}$, and assume $(W_v)_0 = A_v\Sigma_v B_v^T$ is the reduced SVD of $(W_v)_0$, where $A_v \in \mathbb{R}^{D_v\times k_v}$, $B_v \in \mathbb{R}^{d\times k_v}$ and $\Sigma_v \in \mathbb{R}^{k_v\times k_v}$ is diagonal. At each iteration, we calculate $W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1}) \approx \widehat{A}_v\widehat{\Sigma}_v\widehat{B}_v^T$ using the cheaper power method and then filter out the singular vectors $\{\widetilde{A}_v, \widetilde{B}_v\}$ with singular values greater than C_1/η_t. The column and row function subspaces can thus be discovered as the orthonormal bases of $span(A_v, \widetilde{A}_v)$ and $span(B_v, \widetilde{B}_v)$, respectively.

3.3 Optimization Over Combination Weights

Fixing the latent representations and the discovered view function subspaces, the optimization problem w.r.t. the combination weight S_v on the v-th view is

$$\min_{S_v} \ h(S_v) + C_1\|S_v\|_*, \tag{14}$$

where

$$h(W_v) = \frac{1}{nm}\sum_{i=1}^{n}\|x_i^v - A_v S_v B_v^T z_i\|_2^2. \tag{15}$$

Algorithm 1 Streaming View Learning

Input: $\{X^v\}_{v=1}^{m+1}$, $\{z_i\}_{i=1}^n$, $\{A_v, S_v, B_v\}_{v=1}^m$
PART 1- Solve new view function:
Initialize $X_v Z^T (ZZ^T)^{-1} = A_v \Sigma_v B_v^T$, $v = m+1$
for $t = 1, \cdots$, **do**
  $(\widehat{A}_v, \widehat{\Sigma}_v, \widehat{B}_v^T) \leftarrow Power\big(W_v^{t-1} - \frac{1}{\eta_t}\nabla g(W_v^{t-1})\big)$
  $(\widetilde{A}_v, \widetilde{B}_v) \leftarrow SVT(\widehat{A}_v \widehat{\Sigma}_v \widehat{B}_v^T, C_1/\eta_t)$
  $A_v \leftarrow QR([A_v, \widetilde{A}_v])$ and $B_v \leftarrow QR([B_v, \widetilde{B}_v])$
  Solve S_v via Problem (16)
  $(A_v^{'}, S_v^{'}, B_v^{'}) \leftarrow SVD(S_v)$
  $A_v \leftarrow A_v A_v^{'}$, $S_v \leftarrow S_v S_v^{'}$, $B_v \leftarrow B_v B_v^{'}$
  $W_v^t \leftarrow A_v S_v B_v^T$
end for
PART 2- Upgrade view functions:
for $t = 1, \cdots$, **do**
  $\forall\, i$ Solve z_i via Problem (9)
  $\forall\, v$ Solve S_v via Problem (16)
end for
Return $\{A_v, S_v, B_v\}_{v=1}^{m+1}$

Similarly, the proximal gradient technique can be applied to obtain an equivalent objective function,

$$\min_{S_v} \frac{\eta_t}{2}\|S_v - \big(S_v^{t-1} - \frac{1}{\eta_t}\nabla h(S_v^{t-1})\big)\|_F^2 + C_1\|S_v\|_*. \tag{16}$$

Since S_v is a small $k_v \times k_v$ matrix, using the SVT method with an exact SVD operation to solve S_v is feasible.

4. Theoretical Analysis

Here we conduct a theoretical analysis to reveal important properties of the proposed streaming view learning algorithm.

We use the following theorem to show that the latent representations $Z = [z_1, \cdots, z_n]$ become increasingly stable as streaming view learning progresses.

Theorem 1 *Given the latent representations Z^{m-1} learned over $m-1$ views, and Z^m learned over past $m-1$ views and the new m-th view (i.e. m views in total), we have $\|Z^m - Z^{m-1}\|_F = \mathcal{O}(1/m)$.*

Proof Given

$$J_m(Z) = \frac{1}{m}\sum_{v=1}^{m} \ell(X^v, W^v, Z) + C_2\|Z\|_F^2, \tag{17}$$

we have

$$\begin{aligned} J_m(Z)-J_{m-1}(Z) =& \frac{1}{m}\ell(X^m,W^m,Z)+\frac{1}{m}\sum_{v=1}^{m-1}\ell(X^v,W^v,Z) \\ &-\frac{1}{m-1}\sum_{v=1}^{m-1}\ell(X^v,W^v,Z) \\ =&\frac{1}{m}\ell(X^m,W^m,Z)-\frac{1}{m(m-1)}\sum_{v=1}^{m-1}\ell(X^v,W^v,Z). \end{aligned}$$

Since $\ell(\cdot)$ used in the algorithm is Lipschitz in its last argument, $J_m(Z) - J_{m-1}(Z)$ has a Lipschitz constant $\mathcal{O}(1/m)$. Assuming the Lipschitz constant of $J_m(Z) - J_{m-1}(Z)$ is θ_m, we have

$$\begin{aligned} J_{m-1}(Z^m)-J_{m-1}(Z^{m-1}) =& J_{m-1}(Z^m) - J_m(Z^m) + J_m(Z^m) \\ &-J_m(Z^{m-1}) + J_m(Z^{m-1}) - J_{m-1}(Z^m) \\ \leq& J_{m-1}(Z^m) - J_m(Z^m) + J_m(Z^{m-1}) - J_{m-1}(Z^m) \\ \leq& \theta_m\|Z^m - Z^{m-1}\|_F. \end{aligned}$$

Since Z^{m-1} is the minimum of $J_{m-1}(Z)$, we have

$$J_{m-1}(Z^m) - J_{m-1}(Z^{m-1}) \geq 2C_2\|Z^m - Z^{m-1}\|_F^2. \tag{18}$$

Combining the above results, we have,

$$\|Z^m - Z^{m-1}\|_F \leq \frac{\theta_m}{2C_2} = \mathcal{O}(1/m), \tag{19}$$

which completes the proof. ■

Theorem 1 reveals that the streaming views are helpful for deriving a stable multi-view model. We next analyze the optimality of the discovered view function for the newly arrived view using the following theorem.

Theorem 2 *The optimization steps of Part 1 in Algorithm 1 can guarantee that the solved function subspaces of the new view converge to a stationary point.*

The proof of Theorem 2 is listed in supplementary material due to page limitation. This remarkable result shows that the proposed algorithm can efficiently discover the optimal view function subspaces. We next analyze the influence of perturbation of the latent representations on the view function estimation by the following theorem, whose detailed proof is listed in supplementary material.

Theorem 3 *Fixing $\|X^v\| \leq \Upsilon$ for each view. Given latent representation Z with $\|Z\|_F \leq \Omega$, the optimal view function on the v-th view is denoted as W_v. For $\widetilde{Z}$ with $\|\widetilde{Z} - Z\|_F \leq \epsilon$, the optimal view function on the v-th view is defined as $W_v^{'}$. Suppose the smallest eigenvalue of $\widetilde{Z}\widetilde{Z}^T$ is lower bounded by $\lambda > 0$, and both the rank of W_v and $W_v^{'}$ are lower than k. The following error bound holds*

$$\|W_v^{'} - W_v\|_F \leq \frac{1}{\lambda}\big(\Upsilon^2\frac{2\epsilon\Omega+\epsilon^2}{C_1} + \epsilon\Upsilon + 2\sqrt{k+1}\big) \tag{20}$$

This theoretical analysis allows us to summarize as follows. For the newly arrived view, the proposed algorithm is guaranteed to discover its optimal view function based on the convergence analysis in Theorem 2. According to Theorem 1, the learned latent representation will become increasingly stable with more streaming views. Hence, given a small perturbation on latent representation matrix Z, the difference between the target view functions is also small. Most importantly, according to perturbation theory (Li, 1998), it is thus reasonable to assume that the view function subspaces are approximately consistent given the bounded perturbation on the matrix; therefore, it is feasible to simply fine tune the combination weights of these subspaces for better reconstruction. On the other hand, if the number of past views is small, we can use the standard multi-view learning algorithm to re-train the model over past views and the new views together with acceptable resource cost.

5. Experiments

We next evaluated the proposed SVL algorithms for clustering and classification of real-world datasets. The SVL algorithm was compared to canonical correlation analysis (CCA) (Hardoon et al., 2004), the convex multi-view subspace learning algorithm (MCSL) (White et al., 2012), the factorized latent sparse with structured sparsity algorithm (FLSSS) (Jia et al., 2010), and the shared Gaussian process latent variable model (sGPLVM) (Shon et al., 2005). Since these comparison algorithms were not designed for the streaming view setting, we adapted the algorithms for fair comparison such that they employed the idea of multi-view learning to handle new views. Specifically, for each multi-view comparison algorithm, the outputs of the well-trained multi-view model over past views were treated as temporary views, which were then combined with the newly arrived view to train a new multi-view model. Note that we did not adopt the trick to completely re-train multi-view learning algorithms using past views and new views simultaneously, since it is infeasible for practical applications considering the intensive computational cost.

The real-world datasets used in experiments were the *Handwritten Numerals* and *PASCAL VOC'07* datasets. The Handwritten Numerals dataset is composed of $2,000$ data points in 0 to 9 ten-digit classes, where each class contains 200 data points. Six types of features are employed to describe the data: Fourier coefficients of the character shapes (FOU), profile correlations (FAC), Karhunen-Loève coefficients (KAR), pixel averages in 2×3 windows (PIX), Zernike moments (ZER), and morphological features (MOR). The *PASCAL VOC'07* dataset contains around $10,000$ images, each of which was annotated with 20 categories. Sixteen types of features have been used to describe each image including GIST, image tags, 6 color histograms (RGB, LAB, and HSV over single-scale or multi-scale images), and 8 bag-of-features descriptors (SIFT and hue densely extracted or for Harries-Laplacian interest points on single-scale or multi-scale images).

5.1 Multi-view Clustering and Classification

For each algorithm, half of the total views were used for initialization to well train a base multi-view model, and then multi-view learning was conducted with the streaming views. We fixed the dimension of the latent subspace as 100 for different algorithms. Based on the multi-view example subspaces learned through the proposed SVL algorithm and its com-

parison algorithms, the k-means and SVM methods were launched for subsequent clustering and classification, respectively. Clustering performance was assessed by normalized mutual information (NMI) and accuracy (ACC), while classification performance was measured using mean averaged precision (mAP).

Table 1: NMI on the *Handwritten Numerals* dataset. The number following each feature denotes the number of new views processed. 'Single' implies directly launching k-means on the current view.

Algorithm	FAC	FOU	KAR (0)	MOR (1)	PIX (2)	ZER (3)
Single	0.679 ± 0.032	0.547 ± 0.028	0.666 ± 0.030	0.643 ± 0.034	0.703 ± 0.040	0.512 ± 0.025
CCA	-	-	0.755 ± 0.039	0.777 ± 0.038	0.772 ± 0.061	0.799 ± 0.038
FLSSS	-	-	0.833 ± 0.047	0.837 ± 0.029	0.830 ± 0.028	0.840 ± 0.027
sGPLVM	-	-	0.785 ± 0.044	0.794 ± 0.021	0.799 ± 0.056	0.827 ± 0.055
MCSL	-	-	0.798 ± 0.028	0.805 ± 0.047	0.814 ± 0.069	0.815 ± 0.026
SVL	-	-	0.826 ± 0.049	0.840 ± 0.044	0.866 ± 0.037	0.871 ± 0.042

Table 2: ACC on the *Handwritten Numerals* dataset. The number following each feature denotes the number of new views processed. 'Single' implies directly launching k-means on the current view.

Algorithm	FAC	FOU	KAR (0)	MOR (1)	PIX (2)	ZER (3)
Single	0.707 ± 0.065	0.556 ± 0.062	0.689 ± 0.051	0.614 ± 0.058	0.694 ± 0.067	0.534 ± 0.052
CCA	-	-	0.709 ± 0.051	0.710 ± 0.015	0.706 ± 0.037	0.809 ± 0.065
FLSSS	-	-	0.819 ± 0.038	0.831 ± 0.035	0.851 ± 0.027	0.849 ± 0.045
sGPLVM	-	-	0.777 ± 0.052	0.796 ± 0.035	0.788 ± 0.053	0.805 ± 0.058
MCSL	-	-	0.790 ± 0.045	0.804 ± 0.055	0.804 ± 0.036	0.836 ± 0.050
SVL	-	-	0.813 ± 0.052	0.816 ± 0.032	0.908 ± 0.050	0.927 ± 0.050

The performance of different algorithms with respect to the progress of streaming view learning of the clustering task are shown in Tables 1 and 2. Classification results under similar settings are presented in Table 3. In each table, multi-view learning algorithms learn the views from the left to the right column in a streaming manner, such that the results presented on the right side have already been helped by the views on the left. The different multi-view learning algorithms consistently improve clustering and classification performance when more new view information becomes available. Although the base multi-view learning model of the proposed SVL algorithm only achieves comparable or slightly inferior performance to that of comparison algorithms, SVL significantly improves its performance by optimally learning new view functions and upgrading past view functions, such that the advantages of SVL becomes more obvious with increasing of numbers of new views

Table 3: mAP on the *PASCAL VOC'07* dataset. '8' means that multi-view learning begins with 8 views, while the following number denotes the number of new views processed.

Algorithm	8(0)	8(2)	8(4)	8(6)	8(8)
CCA	0.314	0.335	0.347	0.357	0.368
FLSSS	0.507	0.518	0.521	0.532	0.539
sGPLVM	0.441	0.458	0.473	0.487	0.487
MCSL	0.448	0.463	0.469	0.475	0.480
SVL	0.554	0.558	0.562	0.564	0.571

processed. Specifically, in the fifth column of Table 1, the NMI of SVL improves about 20% over that of single-view algorithm and 5% over that of multi-view FLSSS algorithm. On the *PASCAL VOC'07* dataset, training multi-view model over 8 views is already computational expensive, let alone completely re-training with new views.

5.2 Algorithm Analysis

We next varied the dimensionality d of latent representations. The clustering performance of SVL on the *Handwritten Numerals* dataset is presented in Table 4. The performance of the lower-dimensional latent representations is limited, whereas with increased d the latent representations have more power to describe multi-view examples, and the SVL algorithm achieves stable performance.

Table 4: NMI of SVL with different dimensionalities of latent representations on the *Handwritten Numerals* dataset.

d	FAC	FOU	KAR (0)	MOR (1)	PIX (2)	ZER (3)
10	-	-	0.655	0.655	0.688	0.705
20	-	-	0.680	0.762	0.771	0.772
50	-	-	0.755	0.778	0.809	0.838
100	-	-	0.826	0.840	0.866	0.871
150	-	-	0.829	0.834	0.865	0.888

To examine the influence of the number of views used to initialize the base multi-view SVL model, we started SVL with different numbers of views on the *PASCAL VOC'07* dataset. The variability in performance is presented in Figure 1. If the base multi-view model was initialized with insufficient views, the resulting performance was limited (see the first group in Figure 1). This is due to the large estimation error of the functions over the views used to initialize the model. Conversely, if the multi-view model over past views was already well trained, we easily applied SVL to extend the model to handle new views without intensive computational cost, whilst also guaranteeing stable performance improvements. These phenomena are consistent with our theoretical analyses.

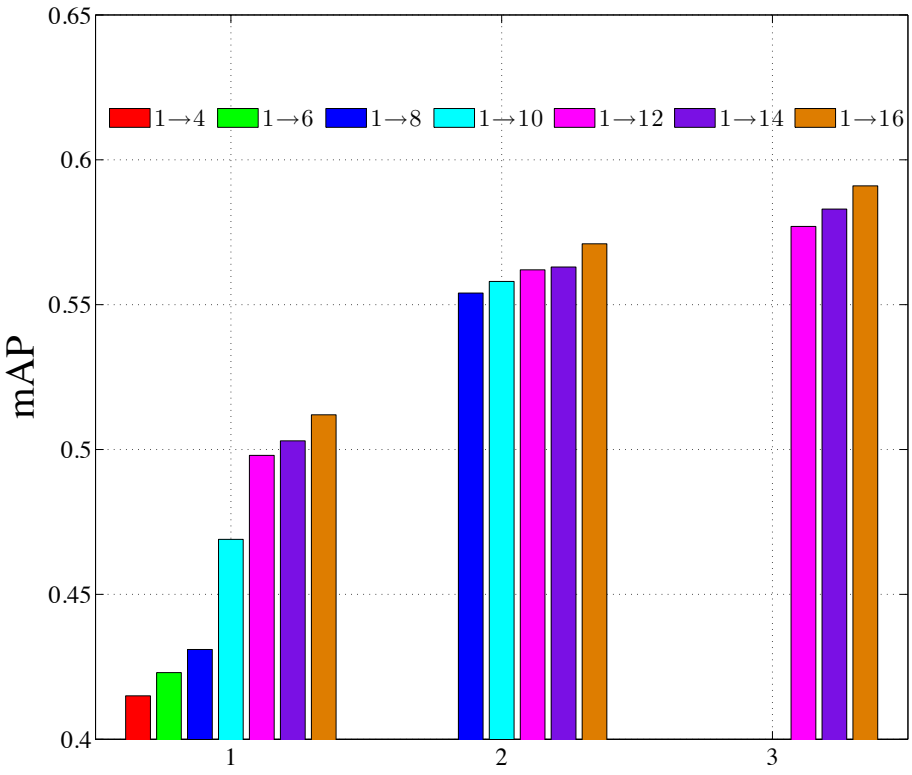

Figure 1: Classification performance of SVL using different numbers of views for initialization on the *PASCAL VOC'07* dataset.

Finally, we examined the influence of the order of streaming views on the learning performance of SVL, and the classification results are shown in Figure 2. For each group in Figure 2, the view order is randomly determined. It can be seen that although classification performance variations is diverse, the resulting performances are roughly equivalent with distinct streaming view orders.

6. Conclusions

Here we investigate the streaming view problem in multi-view learning, in which views arrive in a streaming manner. Instead of discarding the multi-view model trained well over past views, we regard the subspaces of the well-trained view functions as stable and fine-tune the weights for function subspaces combinations while processing new views. In this way, the resulting SVL algorithm can efficiently learn view functions for the new view and update view functions for past views. The convergence issue of the proposed algorithm is theoretically studied, and the influence of streaming views on the multi-view model is addressed. Comprehensive experiments conducted on real-world datasets demonstrate the significance of studying the streaming view problem and the effectiveness of the proposed algorithm.

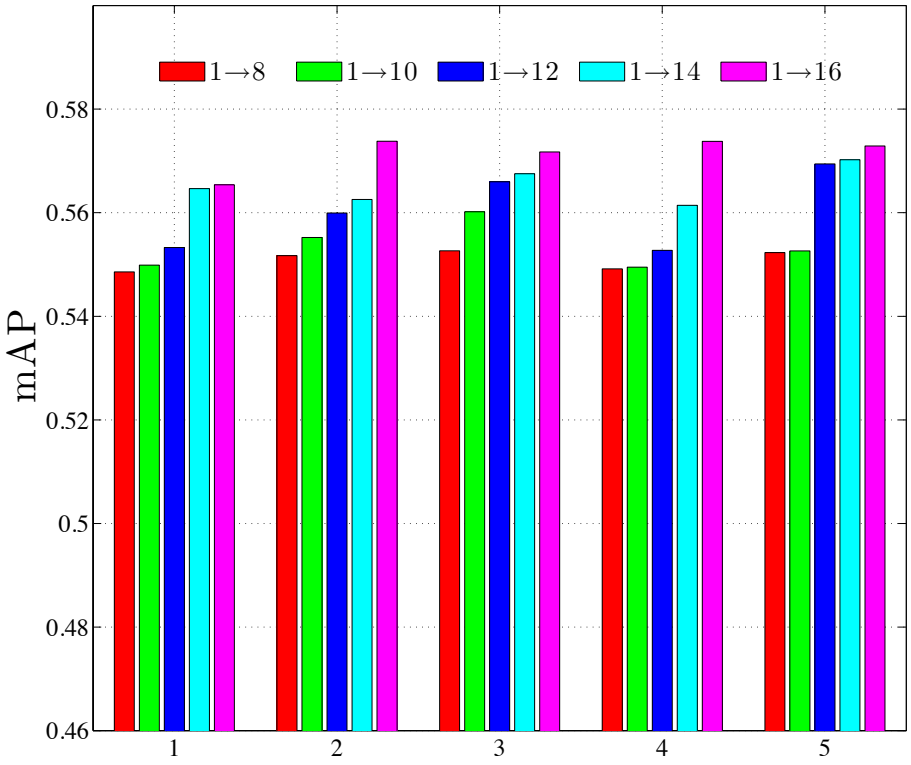

Figure 2: Classification Performance of SVL with distinct view orders on the PASCAL VOC'07 dataset.

References

Peter Arbenz, Daniel Kressner, and D-MATH ETH Zürich. Lecture notes on solving large scale eigenvalue problems. *D-MATH, EHT Zurich*, 2012.

Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100. ACM, 1998.

Jian-Feng Cai, Emmanuel J Candès, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.

Xiao Cai, Feiping Nie, and Heng Huang. Multi-view k-means clustering on big data. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*. Citeseer, 2013.

Kamalika Chaudhuri, Sham M Kakade, Karen Livescu, and Karthik Sridharan. Multi-view clustering via canonical correlation analysis. In *ICML*, 2009.

N. Chen, J. Zhu, F. Sun, and E.P. Xing. Large-margin predictive latent subspace learning for multi-view data analysis. *IEEE Trans. pattern analysis and machine intelligence*, 2012.

Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005.

Navneet Dalal, Bill Triggs, and Cordelia Schmid. Human detection using oriented histograms of flow and appearance. In *Computer Vision–ECCV 2006*, pages 428–441. Springer, 2006.

Osman Güler. *Foundations of optimization*, volume 258. Springer Science & Business Media, 2010.

Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

David Hardoon, Sandor Szedmak, and John Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12):2639–2664, 2004.

Shuiwang Ji and Jieping Ye. An accelerated gradient method for trace norm minimization. In *Proceedings of the 26th annual international conference on machine learning*, pages 457–464. ACM, 2009.

Yangqing Jia, Mathieu Salzmann, and Trevor Darrell. Factorized latent spaces with structured sparsity. In *NIPS*, pages 982–990, 2010.

Xin Jin, Yanzan Zhou, and Bamshad Mobasher. A maximum entropy web recommendation system: combining collaborative and content features. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 612–617. ACM, 2005.

Shaishav Kumar and Raghavendra Udupa. Learning hash functions for cross-view similarity search. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 1360, 2011.

Gert RG Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I Jordan. Learning the kernel matrix with semidefinite programming. *The Journal of Machine Learning Research*, 5:27–72, 2004.

Ivan Laptev, Marcin Marszałek, Cordelia Schmid, and Benjamin Rozenfeld. Learning realistic human actions from movies. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

Ren-Cang Li. Relative perturbation theory: Ii. eigenspace and singular subspace variations. *SIAM Journal on Matrix Analysis and Applications*, 20(2):471–492, 1998.

Aaron Shon, Keith Grochow, Aaron Hertzmann, and Rajesh P Rao. Learning shared latent structure for image synthesis and robotic imitation. In *NIPS*, pages 1233–1240, 2005.

Xiaogang Wang. Intelligent multi-camera video surveillance: A review. *Pattern recognition letters*, 34(1):3–19, 2013.

Martha White, Xinhua Zhang, Dale Schuurmans, and Yao-liang Yu. Convex multi-view subspace learning. In *Advances in Neural Information Processing Systems*, pages 1673–1681, 2012.